



(11) **EP 3 477 596 B1**

(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention
of the grant of the patent:
09.06.2021 Bulletin 2021/23

(51) Int Cl.:
G06T 11/60 ^(2006.01)

(21) Application number: **18197056.7**

(22) Date of filing: **27.09.2018**

(54) **METHOD AND APPARATUS FOR SELECTING A SURFACE IN A LIGHT FIELD, AND
CORRESPONDING COMPUTER PROGRAM PRODUCT**

VERFAHREN UND VORRICHTUNG ZUR AUSWAHL EINER OBERFLÄCHE EINES LICHTFELDS
UND ZUGEHÖRIGES COMPUTERPROGRAMMPRODUKT

PROCÉDÉ ET APPAREIL DE SÉLECTION D'UNE SURFACE DANS UN CHAMP LUMINEUX ET
PRODUIT-PROGRAMME D'ORDINATEUR CORRESPONDANT

(84) Designated Contracting States:
**AL AT BE BG CH CY CZ DE DK EE ES FI FR GB
GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO
PL PT RO RS SE SI SK SM TR**

(30) Priority: **31.10.2017 EP 17306492**

(43) Date of publication of application:
01.05.2019 Bulletin 2019/18

(73) Proprietor: **InterDigital CE Patent Holdings
75017 Paris (FR)**

(72) Inventors:
• **SABATER, Neus
35576 CESSON SÉVIGNÉ (FR)**
• **REINHARD, Erik
35576 CESSON SÉVIGNÉ (FR)**
• **HOG, Matthieu
35576 CESSON SÉVIGNÉ (FR)**
• **GARCES, Elena
35576 CESSON SÉVIGNÉ (FR)**

(74) Representative: **Schmidt-Uhlig, Thomas
Rittner & Partner
Patentanwälte mbB
Schiffgraben 17
30159 Hannover (DE)**

(56) References cited:
US-A1- 2007 236 507 US-A1- 2014 059 462

- **ADRIAN JARABO ET AL: "How do people edit light fields?", ACM TRANSACTIONS ON GRAPHICS (TOG), ACM, US, vol. 33, no. 4, 27 July 2014 (2014-07-27), pages 1-10, XP058051940, ISSN: 0730-0301, DOI: 10.1145/2601097.2601125**
- **IVAN VIOLA ET AL: "Smart Visibility in Visualization", COMPUTATIONAL AESTHETICS IN GRAPHICS, VISUALIZATION AND IMAGING, 18 May 2005 (2005-05-18), pages 209-216, XP058118849, DOI: 10.2312/COMPAESTH/COMPAESTH05/209-216 ISBN: 3-905673-27-4**
- **"3D interaction / Navigating / Global or Local View", BlenderWiki, 15 November 2011 (2011-11-15), XP055437011, Retrieved from the Internet:
URL: https://wiki.blender.org/index.php?title=Doc:2.4/Manual/3D_interaction/Navigating/Global_or_Local_View&oldid=158165
[retrieved on 2017-12-20]**
- **WU GAOCHANG ET AL: "Light Field Image Processing: An Overview", IEEE JOURNAL OF SELECTED TOPICS IN SIGNAL PROCESSING, IEEE, US, vol. 11, no. 7, 1 October 2017 (2017-10-01), pages 926-954, XP011672035, ISSN: 1932-4553, DOI: 10.1109/JSTSP.2017.2747126 [retrieved on 2017-10-25]**

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

EP 3 477 596 B1

Description

1. Field

[0001] The present disclosure relates to light-field imaging, and to technologies for acquiring and processing light-field data. More precisely, the present disclosure generally relates to a method and an apparatus for editing and visualizing a light field based image on a user's device, and finds applications in the domain of image or video rendering.

2. Background

[0002] This section is intended to introduce the reader to various aspects of art, which may be related to various aspects of the present disclosure that are described and/or claimed below. This discussion is believed to be helpful in providing the reader with background information to facilitate a better understanding of the various aspects of the present invention. Accordingly, it should be understood that these statements are to be read in this light, and not as admission of prior art.

[0003] Conventional image capture devices project a three-dimensional scene onto a two-dimensional sensor. During operation, a conventional capture device captures a two-dimensional (2-D) image of the scene representing an amount of light that reaches a photosensor (or photodetector) within the device. However, this 2-D image contains no information about the directional distribution of the light rays that reach the photosensor (which may be referred to as the light-field). Direction of incoming light, for example, is lost during such 2D acquisition and information like depth cannot be recovered for a single system. Thus, a conventional capture device does not store most of the information about the light distribution from the scene.

[0004] Light-field capture devices (also referred to as "light-field data acquisition devices") have been designed to measure a four-dimensional (4D) light-field of the scene by capturing the light from different viewpoints of that scene. Thus, by measuring the amount of light traveling along each beam of light that intersects the photosensor, these devices can capture additional optical information (information about the directional distribution of the bundle of light rays) for providing new imaging applications by post-processing. The information acquired/obtained by a light-field capture device is referred to as the light-field data. Light-field capture devices are defined herein as any devices that are capable of capturing light-field data. There are several types of light-field capture devices, among which:

- plenoptic devices, which use a microlens array placed between the image sensor and the main lens, as described in document US 2013/0222633;
- a camera array.

[0005] The light field data may also be simulated with Computer Generated Imagery (CGI), from a series of 2-D images of a scene each taken from a different viewpoint by the use of a conventional handheld camera.

[0006] Light-field data processing comprises notably, but is not limited to, generating refocused images of a scene, generating perspective views of a scene, generating depth maps of a scene, generating extended depth of field (EDOF) images, generating stereoscopic images, and/or any combination of these.

[0007] Hence, among others, a 4D Light-Field (4DLF) allows computing various re-focused images with adjustable depth-of-field, focalization distances and viewing positions. However, user experience is often limited to simple rendering on TVs or monitors, 2D computers and mobile displays.

[0008] More generally, current light field editing techniques are limited to changing perspective or focus. However, as the number of captured and shared light fields increases, there is an increasing need for editing tools, which would offer the same functions as the well-established editing of 2D images. Actually, image editing programs like Adobe Photoshop® for example provide ways of modifying an object's appearance in a single image by manipulating the pixels of that image.

[0009] Nonetheless, the multidimensional nature of light fields may make common image editing tasks become complex in light field space.

[0010] First, because a light field is a four-dimensional data structure, while most of the existing editing tools and displays are designed for two-dimensional content. Then, because light fields are redundant and associated with multiple views: the editing task performed on one of these multiple views must be propagated in all views for consistency purpose, which is both cumbersome and time-consuming.

[0011] In "Plenoptic image editing", Proceedings of IEEE 6th International Conference on Computer Vision, 1998, Seitz and Kutulakos presented a method of interactive image editing operations designed to maintain consistency between multiple images of a physical 3D scene: edits to any one image propagates automatically to all other images, as if the 3D scene had itself been modified. Hence, the user can quickly modify many images by editing just a few. Propagation to the other images relies on the use of a plenoptic decomposition into separate shape and radiance

components. The propagation mechanism relies on voxel-based reconstruction to obtain pixel correspondence information.

[0012] Such a technique focuses on circular 360° light fields, which are either acquired by rotating the camera around the scene, or by rotating the object to be captured (see figure 4). Hence, such a technique cannot be directly applied to planar light fields, which are either acquired by a camera array or a plenoptic device.

[0013] Moreover, according to this technique, a user needs to navigate through different views of the scene to fully edit the light field, and handle occluded areas: in other words, for some edits which concern occluded areas, the user needs to perform editing on several different images of the scene, in order for it to be propagated to all views of the scene.

[0014] This is both cumbersome and time-consuming.

[0015] It would be desirable to provide a technique for robustly selecting a surface, or material, in a light field that would show improvements over the prior art. Notably, it would be desirable to provide such a technique, which would enable a user to select materials in a light field without the need to navigate through multiple views, in order to edit them.

[0016] In "How do people edit light fields?", ACM Transactions on Graphics, Vol. 33, 2014, pp. 1-10, Jarabo et al. evaluate different light field editing interfaces, tools and workflows from a user perspective. Furthermore, the potential benefits of using depth information when editing as well as the limitations imposed by imperfect depth reconstruction using current techniques are investigated.

[0017] In "Smart visibility in visualization", Computational Aesthetics'05: Proceedings of the First Eurographics conference on Computational Aesthetics in Graphics, Visualization and Imaging, 2005, pp. 209-216, Viola et al. Describe expressive visualization techniques that smartly uncover the most important information in order to maximize the visual information in the resulting images. This is achieved through dynamic changes in visual representations, through deformations, or through spatial modifications of parts of the data.

[0018] The 3D creation suite "**BLENDER**" supports several contextual view modes for its 3D Viewport. These views can change how the overall 3D Viewport looks or how to interact with objects. Global view shows all of the 3D objects in the scene. Local view isolates the selected object or objects, so that they are the only ones visible in the viewport.

1. Summary

[0019] According to an embodiment of the present disclosure, a method for selecting at least one surface in a light field is disclosed. The light field is associated with multiple views corresponding to different view points of a same scene. Such a method comprises:

- displaying one of the multiple views as a reference view;
- selecting at least one surface in the reference view;
- identifying a set of pixels belonging to the at least one surface, which are occluded in the reference view but visible in at least one of the multiple views;
- projecting the set of pixels into the reference view, on the basis of a value of the pixels of the set in the at least one of the multiple views in which they are visible, in order to form an upgraded reference view;
- displaying the upgraded reference view;
- updating said selected surface in said upgraded reference view.

[0020] The present disclosure thus relies on a novel and inventive approach of selection and visualization of materials in light fields. Actually, it allows a user to perform selection of a material, such as a surface or an object, in one reference view, without any need for navigating through all the multiple views associated with the light field to check consistency. To deal with occlusions, the present disclosure relies on projecting the occluded parts in the reference view, so that the user can continue the selection, if needed, in the reference view. Moreover, after visualizing the full surface without occlusion, the user may determine whether all pixels belonging to the surface have been duly selected, or whether the selection should be refined, depending on the type of editing to be performed.

[0021] Such a method is faster and easier than prior art methods, according to which the user needs to navigate through all the views to observe the occluded areas and refine the selection in order to guarantee an accurate selection of the surface or material.

[0022] It must be noted that the light field dataset may be captured by a plurality of image capture devices (such as a camera array) or by a plenoptic image capture device. It may also be simulated with Computer Generated Imagery (CGI).

[0023] According to an embodiment of the present disclosure, such a method also comprises:

- detecting at least one foreground object at least partially occluding the at least one selected surface in the reference view;
- hiding the at least one foreground object in the upgraded reference view.

[0024] Such detection is based on the knowledge of the depth of the points or pixels in the light field, for example based on the associated depth map. If only part of the foreground object occludes the selected surface, this part is hidden.

[0025] According to a further embodiment, hiding the at least one foreground object is performed by applying a transparency effect to the object. Hence, the occluding foreground object (or part of it) may either be totally or partially hidden (e.g. with transparent blending). This may help the user get an overall view of the scene.

[0026] According to yet a further embodiment, such a method also comprises editing the updated selected surface, and the editing belongs to the group comprising:

- changing color;
- changing texture;
- changing illumination;
- scaling;
- stretching.

[0027] More generally, any type of editing commonly performed on a 2D image according to known techniques and existing editing tools may be performed on the surface selected by the user in the light field.

[0028] According to yet a further embodiment, the reference view is a central view of the multiple views.

[0029] This is particularly interesting when the light field acquisition device is a plenoptic camera, such as a Lytro® camera for example, for accuracy purpose: actually, views located on the periphery of the sensor may be subject to optical deformations.

[0030] Embodiments of the present disclosure also concern an apparatus for selecting at least one surface in a light field, the light field being associated with multiple views corresponding to different view points of a same scene. Such an apparatus comprises:

- a display unit for displaying one of the multiple views as a reference view and for displaying an upgraded reference view;
- a selection tool for selecting at least one surface in the reference view and for updating the selected surface in the upgraded reference view;
- a computing module for:
 - identifying a set of pixels belonging to the at least one surface, which are occluded in the reference view but visible in at least one of the multiple views;
 - projecting the set of pixels into the reference view, on the basis of a value of the pixels of the set in the at least one of the multiple views in which they are visible, in order to form the upgraded reference view.

[0031] According to an embodiment of the present disclosure, such an apparatus also comprises an editing module for editing said updated selected surface.

[0032] Such an apparatus may be embedded in a personal computer, a tablet, a smartphone, or any other kind of handheld device.

[0033] All features of the description of the user's device described above in relation to embodiments of the method for selecting a surface in a light field also apply to such a user's device or apparatus.

[0034] The present disclosure also concerns a computer program product downloadable from a communication network and/or recorded on a medium readable by a computer and/or executable by a processor, comprising program code instructions for implementing a method for editing a light field based image on a user's device as described above.

[0035] The present disclosure also concerns a non-transitory computer-readable medium comprising a computer program product recorded thereon and capable of being run by a processor, including program code instructions for implementing a method for editing a light field based image on a user's device as described above.

[0036] Such a computer program may be stored on a computer readable storage medium. A computer readable storage medium as used herein is considered a non-transitory storage medium given the inherent capability to store the information therein as well as the inherent capability to provide retrieval of the information therefrom. A computer readable storage medium can be, for example, but is not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable combination of the foregoing. It is to be appreciated that the following, while providing more specific examples of computer readable storage mediums to which the present principles can be applied, is merely an illustrative and not exhaustive listing as is readily appreciated by one of ordinary skill in the art: a portable computer diskette; a hard disk; a read-only memory (ROM); an erasable programmable read-only memory (EPROM or Flash memory); a portable compact disc read-only memory (CD-ROM); an optical storage device; a magnetic storage device; or any suitable combination of the foregoing.

[0037] It is to be understood that both the foregoing general description and the following detailed description are

exemplary and explanatory and are not restrictive of the invention, as claimed.

[0038] It must also be understood that references in the specification to "one embodiment" or "an embodiment", indicate that the embodiment described may include a particular feature, structure, or characteristic, but every embodiment may not necessarily include the particular feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same embodiment. Further, when a particular feature, structure, or characteristic is described in connection with an embodiment, it is submitted that it is within the knowledge of one skilled in the art to affect such feature, structure, or characteristic in connection with other embodiments whether or not explicitly described.

4. Brief description of the drawings

[0039] Embodiments of the invention can be better understood with reference to the following description and drawings, given by way of example and not limiting the scope of protection, and in which:

- **Figure 1** is a diagram illustrating examples of a plenoptic camera and a multi-array camera;
- **Figure 2** shows a plan view schematically illustrating an example of sensor regions arranged on the sensing surface of the image sensor;
- **Figures 3A-3B-3C** illustrate embodiments of the present disclosure, in which a surface is selected in a reference view of a light-field, for editing purpose;
- **Figure 4** is a flow chart for explaining a process for editing a light-field according to an embodiment of the present disclosure;
- **Figure 5** is a schematic block diagram illustrating an example of an apparatus for editing a light-field according to an embodiment of the present disclosure.

[0040] The components in the figures are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the invention.

5. Detailed description

[0041] The general principle of the present disclosure relies on the use of a single reference view associated with a light field to allow a user to select and visualize a surface, or object, that needs editing. For this purpose, occluded parts of the surface are projected onto this reference view, which allows the user to perform the whole selection of the surface, or material, in the reference view, without having to navigate through all the views to check consistency.

[0042] A description will now be given of a method, an apparatus and a computer-readable storage medium for editing a light field in embodiments of the present disclosure.

[0043] We first recall some general principles applying to light field cameras and images. **Figure 1** is a diagram illustrating examples of a plenoptic camera and a multi-array camera. Light-field cameras are capable of recording 4D light-field data. For instance, 4D light-field cameras may be: a plenoptic camera 200 comprising a main lens 105, a microlens array 110 and an image sensor 115 (**figure 1A**); or a multi-array camera 150 comprising a lens array 155 and a single image sensor 160 (**figure 1B**). For example, a multi-camera array can be a Pelican array camera as the one depicted in the document WO 2014149403 A1.

[0044] In the example of the plenoptic camera 100 as shown in **Figure 1A**, the main lens 105 receives light from an object (not shown) in an object field of the main lens 105 and passes the light through an image field of the main lens 105. The microlens array 110 includes a plurality of microlenses 115 arranged in a two-dimensional array. Obviously, a main lens 105 and microlenses 115 can be more complex optical systems as the ones depicted for example in Figures 12 and 13 of document GB2488905 or document EP2244484 (A1).

[0045] **Figure 2** is a plan view schematically illustrating an example of sensor regions arranged on the sensing surface of the image sensor. As illustrated in **Figure 2**, the image sensor 200 includes a plurality of pixels 205 arranged in a two-dimensional array, and receives the light from the object through the microlens array 210. Each of the microlenses 211 of the microlens array 210 has the lens properties to direct the light to a circular area 215 on the image sensor 200. An outer contour of the circular area 215 may represent the shape of a micro-image formed and captured by a microlens on the image sensor 200, which shape depends on the shape of the microlens 211. Amongst all of the pixels 205 on the image sensor 200, only the pixels 205 located substantially inside the circular area 215 contribute to the imaging. In other words, a pixel area (or sensing area) of pixels 205 that contributes to the imaging is located substantially inside the circular area 215.

[0046] The image sensor 200 of light-field camera records an image comprising a collection of 2D micro-images arranged within a 2D image. Each microlens 211 of the microlens array 210 forms a micro-image represented by a circular area 215. Coordinates of pixels 205 on the sensor 200 is indicated by (x,y) in the x-y coordinate system on the surface of the image sensor 200 as shown in **figure 2**. The distance p shown in **figure 2** is the distance between the

two consecutive micro-images. Microlenses 211 are chosen such that the distance p is larger than the size of the pixel 205. The distance w shown in **figure 2** is the disparity distance between the two consecutive micro-images. Micro-images are referred by their respective coordinate (i,j) in the i - j coordinate system on the surface of the image sensor 200 as shown in **figure 2**.

[0047] As mentioned above, only the pixels 205 located substantially inside the circular area 215 receive the light through the microlens 211. The inter microlens space may be masked out to prevent photons to pass outside a microlens 211 (if the microlenses 211 have a square shape and the inter microlens space is not formed, such a masking is not needed).

[0048] The center of a micro-image (i,j) is located on the image sensor 200 at the coordinate (x_{ij}, y_{ij}) . The θ in **figure 2** represents the angle between the square lattice of pixels 205 and the square lattice of microlenses 211. The coordinate (x_{ij}, y_{ij}) can be deduced by the following equation (1) considering $(x_{0,0}, y_{0,0})$ which is the pixel coordinate of the microlens image (0,0):

$$\begin{bmatrix} x_{i,j} \\ y_{i,j} \end{bmatrix} = p \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} i \\ j \end{bmatrix} + \begin{bmatrix} x_{0,0} \\ y_{0,0} \end{bmatrix} \quad (1)$$

The distances p and w are given in unit of pixel. They are converted into physical unit of distance (meters) P and W , respectively, by multiplying the pixel size δ : $W = \delta w$ and $P = \delta p$. These distances depend on the characteristics of the light field camera.

[0049] In the state of the art, there are several ways to represent (or define) 4D light-field data. Indeed, in the Chapter 3.3 of the Phd dissertation thesis entitled "Digital Light Field Photography" by Ren Ng, published in July 2006, three different ways to represent 4D light-field data are described. Firstly, a 4D light-field data can be represented, when recorded by a plenoptic camera as the one depicted in **Figure 1(a)** for example, by a collection of micro-lens images. 4D light-field data in this representation are named raw images (or raw 4D light-field data). Secondly, 4D light-field data can be represented, either when recorded by a plenoptic camera or by a camera array, by a set of sub-aperture images. A sub-aperture image corresponds to a captured image of a scene from a point of view, the point of view being slightly different between two sub-aperture images. These sub-aperture images give information about the parallax and depth of the imaged scene. Thirdly, 4D light-field data can be represented by a set of epipolar images (see for example the article entitled "Generating EPI Representation of α 4D Light Fields with a Single Lens Focused Plenoptic Camera", by S. Wanner et al., published in the conference proceedings of ISVC 2011).

[0050] In the embodiment of **figure 3A**, the plenoptic capture 31 may be displayed as a 4x4 matrix of views, or sub-aperture images. All sixteen views illustrate the same scene, seen under a different angle, or point of view. According to an embodiment of the present disclosure, one of the views of matrix 31 is selected as a reference view 311. In a preferred embodiment, such a reference view 311 is a central view of the multiple views 31. This is particularly interesting when the light field acquisition device is a plenoptic camera, such as a Lytro® camera for example, for accuracy purpose: actually, views located on the periphery of the sensor may be subject to optical deformations. However, any view of the matrix 31 of views may be selected as a reference view.

[0051] Reference view 311 illustrates a parallelepiped rectangle in the background, which is partially occluded by a ball in the foreground.

[0052] **Figure 3B** illustrates a first exemplary embodiment, in which the user wishes to edit the surface of the ball 32. As illustrated by dashed lines on the left-hand part of **figure 3B**, the user selects the surface 32 on reference view 311. According to embodiments of the present disclosure, the other views of matrix 31 will be used to determine the values of the pixels belonging to the surface of the ball 32, which are self-occluded in reference view 311. The values of these self-occluded pixels are then projected onto reference view 311, in order to form an upgraded reference view 320. On upgraded reference view 320, there are two sets of projected pixels 321 and 322, which values are extracted from extrema views of matrix 31, on which these set of pixels are visible. The projected pixels 321, 322 have a robust match to the other views of the light field where they are visible, in the sense that the mapping is well-defined, i.e. for a point in an extrema view, there is only one corresponding point in the reference view 311.

[0053] **Figure 3C** illustrates a second exemplary embodiment, in which the user wishes to edit the front surface 33 of the parallelepiped rectangle. As illustrated by dashed lines on the left-hand part of **figure 3C**, the user selects the surface 33 on reference view 311. The ball in the foreground partially occludes this surface 33, as may be easily detected by using the depth map: the depth of the selected area 33 d_s is bigger than the depth of the occluding surface d_{occ} . However, there are views in the matrix of views 31, in which pixels belonging to surface 33, which are occluded in reference view 311, are visible. According to embodiments of the present disclosure, the occluding part of the foreground ball is hidden, and the occluded pixels of surface 33 are projected onto reference view 311, using their values in other views of matrix 31 where they are visible, to form an upgraded reference view 330. The projected pixels are denoted as 331 on **figure 3C**. The occluding part of the foreground ball may either be totally erased, or hidden by applying a transparency effect.

[0054] The user may visualize the upgraded reference views 320 and 330, and continue the selection process visualizing the full surface without occlusion (whether self-occlusion in the first embodiment of **figure 3B** or occlusion by a foreground object in the second embodiment of **figure 3C**). He/she may notably check that all desired pixels are indeed selected. Such a selection, and subsequent editing, may be performed by using existing 2D image selection tools, such as Photoshop® for example.

[0055] **Figure 4** shows a flow chart illustrating a process for editing a light-field based image according to an embodiment of the present disclosure.

[0056] Step S401 corresponds to an initialization step, during which the matrix of views 31 of **figure 3A** may be acquired from a light-field acquisition device: a view of the matrix of views 31 is selected as a reference view 311. At step S402, this reference view 311 is displayed on a user's device, such as a personal computer, a tablet or a smartphone for example.

[0057] At step S403, the user selects a surface, or a material, on the displayed reference view 311. The user may perform such a selection by using a mouse, or any other selection tool on the user interface of the user's device. The segmentation of the surface could be pre-computed, for example using methods like the ones disclosed by Hog et al. in "Light Field Segmentation Using a Ray-Based Graph Structure", 2016, European Conference on Computer Vision (Springer International Publishing), or could be done interactively using existing tools for 2D image selection like "Lazy Snapping" or "Magic Wand" from Photoshop®.

[0058] At step S404, pixels belonging to the selected surface, which are not visible in the reference view, but visible in at least one other view of the matrix of views 31 are identified. Such pixels may either be self-occluded pixels, i.e. pixels which are occluded by the surface to which they belong itself, or be occluded by a foreground object.

[0059] At step S405, the identified occluded pixels are projected onto the reference view 311, using their value in the view(s) in which they are visible. The reference view 311 is hence upgraded into a new reference view 320 or 330, in which all occluded parts of the selected surface are visible. In case when some pixels belonging to the surface are occluded by a foreground object, this object may be partially erased or hidden at step S405 as well.

[0060] At step S406, this upgraded reference view 320 or 330 is displayed on the user's device: the user may hence check that all relevant pixels are duly selected, by visualizing the full surface of interest without occlusion. If not, the user may update the surface selection at step S407.

[0061] Thanks to this embodiment of the present disclosure, the user does not need to navigate through all the views of matrix 31 to check that the material selection had been properly done. Actually, the user may robustly select materials in a light field without the need to navigate through the views to fix inconsistencies caused by occlusions or non-Lambertian materials. Moreover, thanks to the projection of the occlusions onto the reference view, existing 2D image selection tools may be used to select the surface of interest on the upgraded reference view. The user is therefore able to coherently select pixels across multiple views in a light field by selecting pixels in only one reference view.

[0062] At step S408, the user may edit the selected surface, for example by changing its color or its glossiness. More generally, thanks to the embodiments of the present disclosure, any 2D image edits that require selection of coherent regions can be applied to the light field associated with the matrix of views 31, thereby significantly expanding the range of edits available for light fields.

[0063] We now explain in greater details a possible embodiment of the present disclosure as regards steps S403 to S405 in **figure 4**, which is based on projection mapping techniques. It must be noted that alternative methods based on patch match optimization could also be used, such as the ones disclosed by Thonat, T., Shechtman, E., Paris, S., & Drettakis, G. (2016) in "Multi-View Inpainting for Image-Based Scene Editing and Rendering", Fourth International Conference on 3D Vision (3DV).

[0064] Assume that a 3D point cloud P is computed from the light field using existing techniques, such as the one disclosed by Sabater, N., Boisson, G., Vandame, B., Kerbiriou, P., Babon, F., Hog, M., et al. (2017) in "Dataset and Pipeline for Multi-view Light-Field Video", Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. It is then possible to define the set of points S that belong to the surface being selected as:

$$S = \{ (x, y, z) | (x, y, z) \in P \}$$

[0065] Among these points, due to the camera location and occlusion of the scene, there are some points, which are visible in the reference view 311 and some points, which are hidden. The reference view is associated with a camera projection function $proj_{\theta}$, which is known a priori, and returns the coordinates of each 3D point in 2D image space:

$$(px, py) = proj_{\theta}(x, y, z).$$

[0066] S_{hidden} are points which belong to the surface S selected at step S403 but which are not visible in the reference view 311:

$$S_{hidden} = \{ (x, y, z) | (x, y, z) \in S \wedge h_{proj\theta}(x, y, z) = 0 \},$$

where $h_{proj\theta}(x, y, z)$ is a visibility function, with value 1 if the point (x, y, z) is visible in the reference view 311, and 0 otherwise. The visibility function $h_{proj\theta}$ depends on the camera projection parameters defined by $proj\theta$.

[0067] Then, it is possible to define another set of points S_{ext} by expanding the set of points S_{hidden} to include points from the visible part of the selected surface S up to a radius, which could be set e.g. to the twenty nearest neighbor points in 3D space.

[0068] With existing techniques, such as the one disclosed by Kazhdan, M. &. (2013) in "Screened poisson surface reconstruction", ACM Transactions on Graphics (TOG), 32(3), 29, a 3D mesh can be reconstructed from the set of points S_{ext} , resulting in a set of vertices: $\{v\}$.

[0069] At step S404, if there is an object occluding the selected surface, the proposed method will hide totally or partially (with transparent blending) the points that belong to such occluding object. These points are found as those with the same 2D image projection as the hidden points, but visible from the reference view.

$$O = \{ (x, y, z) | Proj(x, y, z) = Proj(x', y', z') \wedge (x', y', z') \in S_{hidden} \wedge (x, y, z) \notin S \wedge h_{proj\theta}(x, y, z) = 1 \}$$

Note that these points need not to belong to the same surface being selected.

[0070] At step S405, for projecting the hidden points into the reference view 311, it is built a new 2D planar domain parametrized by $u = (t, s)$, which is local over the area given by the expanded convex hull of the projected points $Proj\theta(S_{ext})$. The amount of expansion of the convex hull might vary depending on image resolution and can be e.g. fifty pixels. Then, the goal is to find a mapping between the 3D mesh $\{v\}$ computed from S_{ext} and the planar domain u subject to certain constraints. These constraints are given by the points which are already visible in the reference view 311 $S_{visible} = \{S_{ext} - S_{hidden}\}$ whose projection is known and fixed $\{u_{fixed}\}$. The mapping is represented by a vector $u(v) = [t(v); s(v)]$. The mapping is parameterized by a uniform grid, where each cell of the grid is a quadrilateral. For each point of the mesh $\{v\}$, the value of its corresponding point u in the output domain is computed via the following optimization:

$$\min_{u, T} \sum_{i=1}^M \lambda_1 |u_i - u_{fixed}|^2 + \lambda_2 |u_i - T_i(v_i)|^2 + \lambda_3 |\mathcal{L}(u_i) - \mathcal{L}(v_i)|^2$$

where the first term represents the point constraint, corresponding to the points already visible in the reference view 311 that will serve as anchor points. The second term is the projection constraint, where T_i is the transformation function that it is needed to apply to the vertex of the mesh. The third term, is the structural constraint, and it guarantees the smoothness and continuity of the result. In this case, $\mathcal{L}$ is the Laplacian function per vertex as defined by Sorkine, O., Cohen-Or, D., Lipman, Y., Alexa, M., Rössl, C., & Seidel, H. P. (2004) in "Laplacian Surface Editing", Eurographics/ACM SIGGRAPH symposium on Geometry processing, but alternatives could be used (see for example Igarashi, T. I. (2009) "Implementing As-Rigid-As-Possible Shape Manipulation and Surface Flattening", Journal of Graphics, GPU, and Game Tools). Both u and T can be solved by linear least squares optimization, for example, following the approach disclosed by Igarashi in the above document. The constants λ_1 , λ_2 , and λ_3 control the degree of influence of each term and can be tuned by the user. M is the number of points of the mesh $\{v\}$.

[0071] Figure 5 is a schematic block diagram illustrating an example of an apparatus for editing a light field based image, according to an embodiment of the present disclosure.

[0072] An apparatus 50 illustrated in Figure 5 includes a processor 51, a storage unit 52, an input device 53, an output device 54, and an interface unit 55 which are connected by a bus 56. Of course, constituent elements of the computer apparatus 50 may be connected by a connection other than a bus connection using the bus 56. Such an apparatus 50 is for example a Personal Computer, a tablet, or any other type of handheld device, such as a smartphone.

[0073] The processor 51 controls operations of the apparatus 50. The storage unit 52 stores at least one program to be executed by the processor 51, and various data, including data of 4D light field images captured and provided by a light field camera, parameters used by computations performed by the processor 51, intermediate data of computations performed by the processor 51, and so on. The processor 51 may be formed by any known and suitable hardware, or software, or a combination of hardware and software. For example, the processor 51 may be formed by dedicated hardware such as a processing circuit, or by a programmable processing unit such as a CPU (Central Processing Unit) that executes a program stored in a memory thereof.

[0074] The storage unit 52 may be formed by any suitable storage or means capable of storing the program, data, or the like in a computer-readable manner. Examples of the storage unit 52 include non-transitory computer-readable storage media such as semiconductor memory devices, and magnetic, optical, or magneto-optical recording media loaded into a read and write unit. The program causes the processor 51 to perform a process for editing a light field according to an embodiment of the present disclosure as described above with reference to **Figure 4**.

[0075] The input device 53 may be formed by a keyboard, a pointing device such as a mouse, or the like for use by the user to input commands and select surfaces or objects on a reference view. The output device 54 may be formed by a display device to display, for example, a Graphical User Interface (GUI), or a reference view of the light field, or an upgraded reference view, onto which occluded pixels have been projected, according to embodiments of the present disclosure. The input device 53 and the output device 54 may be formed integrally by a touchscreen panel, for example.

[0076] The interface unit 55 provides an interface between the apparatus 50 and an external apparatus. The interface unit 55 may be communicable with the external apparatus via cable or wireless communication. In this embodiment, the external apparatus may be a light field camera. In this case, data of 4D light field images captured by the light field camera can be input from the light field camera to the apparatus 50 through the interface unit 55, then stored in the storage unit 52, and displayed on screen 54.

[0077] The apparatus 50 and the light field camera may communicate with each other via cable or wireless communication.

[0078] Although only one processor 51 is shown on **figure 5**, it must be understood that such a processor may comprise different modules and units embodying the functions carried out by apparatus 50 according to embodiments of the present disclosure, such as:

- a selection tool for selecting a surface in the reference view displayed on output device 54 and for updating the selected surface in the upgraded reference view displayed on output device 54;
- a computing module for:
 - identifying a set of pixels belonging to the selected surface, which are occluded in the reference view but visible in another view;
 - projecting the set of pixels into the reference view, on the basis of a value of these pixels in the other view in which they are visible, in order to form the upgraded reference view;
- an editing module for editing the updated selected surface.

[0079] These modules and units may also be embodied in several processors 51 communicating and co-operating with each other.

[0080] As will be appreciated by one skilled in the art, aspects of the present principles can be embodied as a system, method or computer readable medium. Accordingly, aspects of the present principles can take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, and so forth), or an embodiment combining software and hardware aspects that can all generally be referred to herein as a "circuit", "module", or "system".

[0081] When the present principles are implemented by one or several hardware components, it can be noted that a hardware component comprises a processor that is an integrated circuit such as a central processing unit, and/or a microprocessor, and/or an Application-specific integrated circuit (ASIC), and/or an Application-specific instruction-set processor (ASIP), and/or a graphics processing unit (GPU), and/or a physics processing unit (PPU), and/or a digital signal processor (DSP), and/or an image processor, and/or a coprocessor, and/or a floating-point unit, and/or a network processor, and/or an audio processor, and/or a multi-core processor. Moreover, the hardware component can also comprise a baseband processor (comprising for example memory units, and a firmware) and/or radio electronic circuits (that can comprise antennas) which receive or transmit radio signals. In one embodiment, the hardware component is compliant with one or more standards such as ISO/IEC 18092 / ECMA-340, ISO/IEC 21481 / ECMA-352, GSM, STOLPaN, ETSI/SCP (Smart Card Platform), GlobalPlatform (i.e. a secure element). In a variant, the hardware component is a Radio-frequency identification (RFID) tag. In one embodiment, a hardware component comprises circuits that enable Bluetooth communications, and/or Wi-fi communications, and/or Zigbee communications, and/or USB communications and/or Firewire communications and/or NFC (for Near Field) communications.

[0082] Furthermore, aspects of the present principles can take the form of a computer readable storage medium. Any combination of one or more computer readable storage medium(s) may be utilized.

[0083] Thus for example, it will be appreciated by those skilled in the art that the block diagrams presented herein represent conceptual views of illustrative system components and/or circuitry embodying the principles of the invention. Similarly, it will be appreciated that any flow charts, flow diagrams, state transition diagrams, pseudo code, and the like represent various processes which may be substantially represented in computer readable storage media and so ex-

cuted by a computer or a processor, whether or not such computer or processor is explicitly shown.

[0084] Although the present disclosure has been described with reference to one or more examples, workers skilled in the art will recognize that changes may be made in form and detail.

Claims

1. A method for selecting at least one surface in a light field, said light field being associated with multiple views (31) corresponding to different view points of a same scene, the method comprising:

- displaying (S402) one of said multiple views (31) as a reference view (311); and
- selecting (S403) at least one surface (32, 33) in said reference view (311); **characterized in that** the method further comprises:
 - identifying (S404) a set of pixels belonging to said at least one surface (32, 33), which are occluded in said reference view (311) but visible in at least one of said multiple views (31);
 - projecting (S405) said set of pixels into said reference view (311), on the basis of a value of said pixels of said set in said at least one of said multiple views (31) in which they are visible, in order to form an upgraded reference view (320, 330);
 - displaying (S406) said upgraded reference view (320, 330); and
 - updating (S407) said selected surface (32, 33) in said upgraded reference view (320, 330).

2. The method of claim 1, wherein further comprising:

- detecting at least one foreground object (32) at least partially occluding said at least one surface (33) in said reference view (311); and
- hiding said at least one foreground object (32) in said upgraded reference view (330).

3. The method of claim 2, wherein hiding said at least one foreground object (32) is performed by applying a transparency effect to said object (32).

4. The method of any of claims 1 to 3, further comprising editing (S408) said updated selected surface (32, 33), wherein said editing (S408) belongs to the group comprising:

- changing color;
- changing texture;
- changing illumination;
- scaling; and
- stretching.

5. The method of any of the previous claims, wherein said reference view (311) is a central view of said multiple views (31).

6. An apparatus (50) for selecting at least one surface in a light field, said light field being associated with multiple views (31) corresponding to different view points of a same scene, the apparatus (50) comprising:

- a display unit (54) for displaying (S402) one of said multiple views (31) as a reference view (311) and for displaying (S406) an upgraded reference view (320, 330);
- a selection tool (53) for selecting (S403) at least one surface (32, 33) in said reference view (311) and for updating (S407) said selected surface (32, 33) in said upgraded reference view (320, 330); and
- a computing module (51) for:

- identifying (S404) a set of pixels belonging to said at least one surface (32, 33), which are occluded in said reference view (311) but visible in at least one of said multiple views (31); and
- projecting (S405) said set of pixels into said reference view (311), on the basis of a value of said pixels of said set in said at least one of said multiple views (31) in which they are visible, in order to form said upgraded reference view (320, 330).

7. The apparatus of claim 6, further comprising an editing module for editing (S408) said updated selected surface

(32, 33).

8. The apparatus of claim 6, wherein the computing module (51) is further configured for:

- detecting at least one foreground object (32) at least partially occluding said at least one surface (33) in said reference view (311);
- hiding said at least one foreground object (32) in said upgraded reference view (330).

9. The apparatus of claim 8, wherein hiding said at least one foreground object (32) is performed by applying a transparency effect to said object (32).

10. A computer program product downloadable from a communication network and/or recorded on a medium readable by a computer and/or executable by a processor, comprising program code instructions for implementing a method according to any of claims 1 to 5.

11. A non-transitory computer-readable medium comprising a computer program product recorded thereon and capable of being run by a processor, including program code instructions for implementing a method according to any of claims 1 to 5.

Patentansprüche

1. Verfahren zum Auswählen wenigstens einer Oberfläche in einem Lichtfeld, wobei das Lichtfeld mit mehreren Ansichten (31) verknüpft ist, die verschiedenen Perspektiven einer selben Szene entsprechen, wobei das Verfahren umfasst:

- Anzeigen (S402) einer der mehreren Ansichten (31) als eine Referenzansicht (311); und
- Auswählen (S403) wenigstens einer Oberfläche (32, 33) in der Referenzansicht (311);

dadurch gekennzeichnet, dass das Verfahren ferner umfasst:

- Identifizieren (S404) einer Menge von Pixeln, die zu der wenigstens einen Oberfläche (32, 33) gehören, die in der Referenzansicht (311) verdeckt sind, die aber in wenigstens einer der mehreren Ansichten (31) sichtbar sind;
- Projizieren (S405) der Menge von Pixeln in die Referenzansicht (311) auf der Grundlage eines Werts der Pixel der Menge in der wenigstens einen der mehreren Ansichten (31), in denen sie sichtbar sind, um eine erweiterte Referenzansicht (320, 330) zu bilden;
- Anzeigen (S406) der erweiterten Referenzansicht (320, 330); und
- Aktualisieren (S407) der ausgewählten Oberfläche (32, 33) in der aktualisierten Referenzansicht (320, 330).

2. Verfahren nach Anspruch 1, das ferner umfasst:

- Detektieren wenigstens eines Vordergrundobjekts (32), das die wenigstens eine Oberfläche (33) in der Referenzansicht (311) wenigstens teilweise verdeckt; und
- Verbergen des wenigstens einen Vordergrundobjekts (32) in der erweiterten Referenzansicht (330).

3. Verfahren nach Anspruch 2, wobei das Verbergen des wenigstens einen Vordergrundobjekts (32) durch Anwenden eines Transparenzeffekts auf das Objekt (32) ausgeführt wird.

4. Verfahren nach einem der Ansprüche 1 bis 3, das ferner das Bearbeiten (S408) der aktualisierten ausgewählten Oberfläche (32, 33) umfasst, wobei das Bearbeiten (S408) zu der Gruppe gehört, die Folgendes umfasst:

- Ändern der Farbe;
- Ändern der Textur;
- Ändern der Beleuchtung;
- Skalieren; und
- Dehnen.

5. Verfahren nach einem der vorhergehenden Ansprüche, wobei die Referenzansicht (311) eine zentrale Ansicht der mehreren Ansichten (31) ist.

6. Vorrichtung (50) zum Auswählen wenigstens einer Oberfläche in einem Lichtfeld, wobei das Lichtfeld mit mehreren Ansichten (31) verknüpft ist, die verschiedenen Perspektiven einer selben Szene entsprechen, wobei die Vorrichtung (50) umfasst:

- eine Anzeigeeinheit (54) zum Anzeigen (S402) einer der mehreren Ansichten (31) als eine Referenzansicht (311) und zum Anzeigen (S406) einer erweiterten Referenzansicht (320, 330);
- ein Auswahlhilfsmittel (53) zum Auswählen (S403) wenigstens einer Oberfläche (32, 33) in der Referenzansicht (311) und zum Aktualisieren (S407) der ausgewählten Oberfläche (32, 33) in der erweiterten Referenzansicht (320, 330); und
- ein Berechnungsmodul (51) zum:

- Identifizieren (S404) einer Menge von Pixeln, die zu der wenigstens einen Oberfläche (32, 33) gehören, die in der Referenzansicht (311) verdeckt sind, die aber in wenigstens einer der mehreren Ansichten (31) sichtbar sind; und
- Projizieren (S405) der Menge von Pixeln in die Referenzansicht (311) auf der Grundlage eines Werts der Pixel der Menge in der wenigstens einen der mehreren Ansichten (31), in denen sie sichtbar sind, um die erweiterte Referenzansicht (320, 330) zu bilden.

7. Vorrichtung nach Anspruch 6, die ferner ein Bearbeitungsmodul zum Bearbeiten (S408) der aktualisierten ausgewählten Oberfläche (32, 33) umfasst.

8. Vorrichtung nach Anspruch 6, wobei das Berechnungsmodul (51) ferner konfiguriert ist zum:

- Detektieren wenigstens eines Vordergrundobjekts (32), das die wenigstens eine Oberfläche (33) in der Referenzansicht (311) wenigstens teilweise verdeckt;
- Verbergen des wenigstens einen Vordergrundobjekts (32) in der erweiterten Referenzansicht (330).

9. Vorrichtung nach Anspruch 8, wobei das Verbergen des wenigstens einen Vordergrundobjekts (32) durch Anwenden eines Transparenzeffekts auf das Objekt (32) ausgeführt wird.

10. Computerprogrammprodukt, das von einem Kommunikationsnetz heruntergeladen werden kann und/oder das auf einem durch einen Computer lesbaren Medium aufgezeichnet ist und/oder das durch einen Prozessor ausführbar ist, umfassend Programmcodeanweisungen zum Implementieren eines Verfahrens nach einem der Ansprüche 1 bis 5.

11. Nichttransitorisches computerlesbares Medium, das ein Computerprogrammprodukt umfasst, das darauf aufgezeichnet ist und durch einen Prozessor ausgeführt werden kann, enthaltend Programmcodeanweisungen zum Implementieren eines Verfahrens nach einem der Ansprüche 1 bis 5.

Revendications

1. Un procédé pour la sélection au moins d'une surface dans un champ lumineux, ledit champ lumineux étant associé à des vues multiples (31) correspondant à des points de vue différents d'une même scène, le procédé comprenant :

- l'affichage (S402) de l'une desdites vues multiples (31) comme vue de référence (311); et
- la sélection (S403) au moins d'une surface (32, 33) dans ladite vue de référence (311);

caractérisé en ce que le procédé comprend en outre :

- l'identification (S404) d'un ensemble de pixels faisant partie de ladite au moins une surface (32, 33), lesquels sont occultés dans ladite vue de référence (311), mais visibles dans au moins l'une desdites vues multiples (31);
- la projection (S405) dudit ensemble de pixels dans ladite vue de référence (311), sur la base d'une valeur desdits pixels dudit ensemble dans ladite au moins une desdites vues multiples (31) dans laquelle ils sont visibles, afin de former une vue de référence améliorée (320, 330);

- l'affichage (S406) de ladite vue de référence améliorée (320, 330); et
- la mise à jour (S407) de ladite surface sélectionnée (32, 33) dans ladite vue de référence actualisée (320, 330).

2. Le procédé de la revendication 1, comprenant en outre :

- la détection au moins d'un objet de premier plan (32) occultant au moins partiellement ladite au moins une surface (33) dans ladite vue de référence (311); et
- le masquage dudit au moins un objet de premier plan (32) dans ladite vue de référence améliorée (330).

3. Le procédé de la revendication 2, dans lequel le masquage dudit au moins un objet de premier plan (32) est réalisé en appliquant un effet de transparence sur ledit objet (32).

4. Le procédé de l'une des revendications 1 à 3, comprenant en outre l'édition (S408) de ladite surface sélectionnée mise à jour (32, 33), dans lequel ladite édition (S408) fait partie du groupe comprenant :

- le changement de couleur;
- le changement de texture;
- le changement d'illumination;
- la graduation; et
- l'allongement.

5. Le procédé de l'une des revendications précédentes, dans lequel ladite vue de référence (311) est une vue centrale desdites vues multiples (31).

6. Un appareil (50) pour la sélection au moins d'une surface dans un champ lumineux, ledit champ lumineux étant associé à des vues multiples (31) correspondant à des points de vue différents d'une même scène, l'appareil (50) comprenant :

- une unité d'affichage (54) pour l'affichage (S402) de l'une desdites vues multiples (31) comme vue de référence (311) et pour l'affichage (S406) d'une vue de référence améliorée (320, 330);
- un outil de sélection (53) pour la sélection (S403) au moins d'une surface (32, 33) dans ladite vue de référence (311) et pour la mise à jour (S407) de ladite surface sélectionnée (32, 33) dans ladite vue de référence améliorée (320, 330); et
- un module informatique (51) pour :

- l'identification (S404) d'un ensemble de pixels faisant partie de ladite au moins une surface (32, 33), lesquels sont occultés dans ladite vue de référence (311), mais visibles dans au moins l'une desdites vues multiples (31); et
- la projection (S405) dudit ensemble de pixels dans ladite vue de référence (311), sur la base d'une valeur desdits pixels dudit ensemble dans ladite au moins une desdites vues multiples (31) dans laquelle ils sont visibles, afin de former ladite vue de référence améliorée (320, 330).

7. L'appareil de la revendication 6, comprenant en outre un module d'édition pour l'édition (S408) de ladite surface sélectionnée actualisée (32, 33).

8. L'appareil de la revendication 6, dans lequel le module informatique (51) est en outre configuré pour :

- la détection au moins d'un objet de premier plan (32) occultant au moins partiellement ladite au moins une surface (33) dans ladite vue de référence (311);
- le masquage dudit au moins un objet de premier plan (32) dans ladite vue de référence améliorée (330).

9. L'appareil de la revendication 8, dans lequel le masquage dudit au moins un objet de premier plan (32) est réalisé en appliquant un effet de transparence sur ledit objet (32).

10. Un produit de programme informatique téléchargeable depuis un réseau de communication et/ou enregistré sur un support lisible par un ordinateur et/ou exécutable par un processeur, comprenant des instructions de code de programme pour la mise en œuvre d'un procédé conformément à l'une des revendications 1 à 5.

EP 3 477 596 B1

11. Un support lisible sur ordinateur non-transitoire, comprenant un produit de programme informatique enregistré sur celui-ci et pouvant être exécuté par un processeur, incluant des instructions de code de programme pour la mise en œuvre d'un procédé conformément à l'une des revendications 1 à 5.

5

10

15

20

25

30

35

40

45

50

55




Figure 1




Figure 2




Figure 5




Fig. 3A




Fig. 3B





Fig. 3C

Figure 4



REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US 20130222633 A [0004]
- WO 2014149403 A1 [0043]
- GB 2488905 A [0044]
- EP 2244484 A1 [0044]

Non-patent literature cited in the description

- Plenoptic image editing. *Proceedings of IEEE 6th International Conference on Computer Vision*, 1998 [0011]
- JARABO. How do people edit light fields?. *ACM Transactions on Graphics*, 2014, vol. 33, 1-10 [0016]
- VIOLA. Smart visibility in visualization. *Computational Aesthetics'05: Proceedings of the First Eurographics conference on Computational Aesthetics in Graphics, Visualization and Imaging*, 2005, 209-216 [0017]
- Digital Light Field Photography. REN NG. Phd dissertation thesis. July 2006 [0049]
- S. WANNER et al. Generating EPI Representation of α 4D Light Fields with a Single Lens Focused Plenoptic Camera. *conference proceedings of ISVC*, 2011 [0049]
- Light Field Segmentation Using a Ray-Based Graph Structure. HOG et al. *European Conference on Computer Vision*. Springer International Publishing, 2016 [0057]
- THONAT, T. ; SHECHTMAN, E. ; PARIS, S. ; DRET-TAKIS, G. Multi-View Inpainting for Image-Based Scene Editing and Rendering. *Fourth International Conference on 3D Vision (3DV)*, 2016 [0063]
- SABATER, N. ; BOISSON, G. ; VANDAME, B. ; KERBIRIOU, P. ; BABON, F. ; HOG, M. et al. Dataset and Pipeline for Multi-view Light-Field Video. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017 [0064]
- KAZHDAN, M. Screened poisson surface reconstruction. *ACM Transactions on Graphics (TOG)*, 2013, vol. 32 (3), 29 [0068]
- SORKINE, O. ; COHEN-OR, D. ; LIPMAN, Y. ; ALEXA, M. ; RÖSSL, C. ; SEIDEL, H. P. Laplacian Surface Editing. *Eurographics/ACM SIGGRAPH symposium on Geometry processing*, 2004 [0070]
- IGARASHI, T. I. Implementing As-Rigid-As-Possible Shape Manipulation and Surface Flattening. *Journal of Graphics, GPU, and Game Tools*, 2009 [0070]